

Ref #	Hits	Search Query	DBs	Default Operator	Plurals	Time Stamp
L1	4	(Radek near Aster).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:18
L2	14	(susan near coatney).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:19
L3	7636	network adj (cache or buffer or storage)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:19
L4	1041	(network or proxy) adj cache	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:19
L5	353244	server	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:19
L6	2224	port adj (id or identification)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:19
L7	2531698	switch\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L8	815313	disk or disks	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L9	235	ID near router	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L10	331573	internet	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20

L11	61799	router	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L12	9661	(fiber or fibre) adj channel	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L13	1	(disk adj (ID or identification)) same (offline)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L14	14685	offline	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L15	22476	(storage\$2 or unit\$2 or device\$2 or disk\$2 or server\$2) adj (id or identification)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L16	7636	network adj (cache or buffer or storage)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L17	1041	(network or proxy) adj cache	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L18	510	L16 and L17	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L19	2224	port adj (id or identification)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L20	510	L16 and L17	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20

L21	1	L20 and L19	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L22	2531698	switch\$4	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L23	206	L20 and L22	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L24	22476	(storage\$2 or unit\$2 or device\$2 or disk\$2 or server\$2) adj (id or identification)	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L25	206	L20 and L22	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L26	10	L25 and L24	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L27	14685	offline	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L28	10	L25 and L24	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L29	0	L28 and L27	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
L30	14	(susan near coatney).in.	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20

L31	0	L30 and L28	US-PGPUB; USPAT; EPO; JPO; DERWENT; IBM_TDB	OR	OFF	2005/10/28 13:20
-----	---	-------------	---	----	-----	------------------


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

 Search: ☒ The ACM Digital Library ☐ The Guide


[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used

[network](#) [proxy](#) [cache](#) [buffer](#) [queue](#) [fibre channel](#) [offline](#)

Found 208 of 166,357

Sort results by

 ☒
[Save results to a Binder](#)[Try an Advanced Search](#)[Try this search in The ACM Guide](#)

Display results

 ☒
[Search Tips](#)
☐ Open results in a new window

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

Best 200 shown

Relevance scale ☐ ☐ ☐ ☐ ☐

1 [Empirical evaluation of multi-level buffer cache collaboration for storage systems](#)



Zhifeng Chen, Yan Zhang, Yuanyuan Zhou, Heidi Scott, Berni Schiefer

 June 2005 **ACM SIGMETRICS Performance Evaluation Review , Proceedings of the 2005 ACM SIGMETRICS international conference on Measurement and modeling of computer systems SIGMETRICS '05**, Volume 33 Issue 1

Publisher: ACM Press

Full text available: [pdf\(379.25 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

To bridge the increasing processor-disk performance gap, buffer caches are used in both storage clients (e.g. database systems) and storage servers to reduce the number of slow disk accesses. These buffer caches need to be managed effectively to deliver the performance commensurate to the aggregate buffer cache size. To address this problem, two paradigms have been proposed recently to *collaboratively* manage these buffer caches together: the **hierarchy-aware caching** maintains ...

Keywords: collaborative caching, database, file system, storage system

2 [TinyDB: an acquisitional query processing system for sensor networks](#)



Samuel R. Madden, Michael J. Franklin, Joseph M. Hellerstein, Wei Hong

 March 2005 **ACM Transactions on Database Systems (TODS)**, Volume 30 Issue 1

Publisher: ACM Press

Full text available: [pdf\(1.67 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We discuss the design of an acquisitional query processor for data collection in sensor networks. Acquisitional issues are those that pertain to where, when, and how often data is physically acquired (*sampled*) and delivered to query processing operators. By focusing on the locations and costs of acquiring data, we are able to significantly reduce power consumption over traditional passive systems that assume the a priori existence of data. We discuss simple extensions to SQL for controll...

Keywords: Query processing, data acquisition, sensor networks


3 [Building reliable mobile-aware applications using the Rover toolkit](#)



Anthony D. Joseph, M. Frans Kaashoek

 October 1997 **Wireless Networks**, Volume 3 Issue 5

Publisher: Kluwer Academic Publishers

Full text available:  pdf(371.04 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper discusses extensions to the Rover toolkit for constructing reliable mobile-aware applications. The extensions improve upon the existing failure model, which addresses client or communication failures and guarantees reliable message delivery from clients to server, but does not address server failures (e.g., the loss of an incoming message due to server failure) (Joseph et al., 1997). Due to the unpredictable, intermittent communication connectivity typically found in mobile client ...

4 The click modular router



Eddie Kohler, Robert Morris, Benjie Chen, John Jannotti, M. Frans Kaashoek
August 2000 **ACM Transactions on Computer Systems (TOCS)**, Volume 18 Issue 3

Publisher: ACM Press

Full text available:  pdf(376.31 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Clicks is a new software architecture for building flexible and configurable routers. A Click router is assembled from packet processing modules called elements. Individual elements implement simple router functions like packet classification, queuing, scheduling, and interfacing with network devices. A router configurable is a directed graph with elements at the vertices; packets flow along the edges of the graph. Several features make individual elements more powerful and ...

Keywords: component systems, routers, software router performance

5 Building reliable mobile-aware applications using the Rover toolkit



Anthony D. Joseph, M. Frans Kaashoek
November 1996 **Proceedings of the 2nd annual international conference on Mobile computing and networking**

Publisher: ACM Press

Full text available:  pdf(1.36 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

6 DHTTP: an efficient and cache-friendly transfer protocol for the web



Michael Rabinovich, Hua Wang
December 2004 **IEEE/ACM Transactions on Networking (TON)**, Volume 12 Issue 6

Publisher: IEEE Press

Full text available:  pdf(487.71 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Today's HTTP carries Web interactions over client-initiated TCP connections. An important implication of using this transport method is that interception caches in the network violate the end-to-end principle of the Internet, which severely limits deployment options of these caches. Furthermore, while an increasing number of Web interactions are short, and in fact frequently carry only control information and no data, TCP is often inefficient for short interactions. We propose a new transfer protocol ...

Keywords: HTTP protocol, interception caching, internet, web performance

7 Run-time adaptation in river



Remzi H. Arpaci-Dusseau
February 2003 **ACM Transactions on Computer Systems (TOCS)**, Volume 21 Issue 1

Publisher: ACM Press

Full text available:  [pdf\(849.04 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present the design, implementation, and evaluation of run-time adaptation within the River dataflow programming environment. The goal of the River system is to provide adaptive mechanisms that allow database query-processing applications to cope with performance variations that are common in cluster platforms. We describe the system and its basic mechanisms, and carefully evaluate those mechanisms and their effectiveness. In our analysis, we answer four previously unanswered and important que ...

Keywords: Performance availability, clusters, parallel I/O, performance faults, robust performance, run-time adaptation

8 Robustness: Defensive programming: using an annotation toolkit to build DoS-resistant software



Xiaohu Qie, Ruoming Pang, Larry Peterson

December 2002 **ACM SIGOPS Operating Systems Review**, Volume 36 Issue SI

Publisher: ACM Press

Full text available:  [pdf\(2.13 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

This paper describes a toolkit to help improve the robustness of code against DoS attacks. We observe that when developing software, programmers primarily focus on functionality. Protecting code from attacks is often considered the responsibility of the OS, firewalls and intrusion detection systems. As a result, many DoS vulnerabilities are not discovered until the system is attacked and the damage is done. Instead of reacting to attacks after the fact, this paper argues that a better solution i ...

9 Experimental testbeds and data: Performance optimizations for wireless wide-area networks: comparative study and experimental evaluation



Rajiv Chakravorty, Suman Banerjee, Pablo Rodriguez, Julian Chesterfield, Ian Pratt

September 2004 **Proceedings of the 10th annual international conference on Mobile computing and networking**

Publisher: ACM Press

Full text available:  [pdf\(262.46 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

We present a comparative performance study of a wide selection of optimization techniques to enhance application performance in the context of wide-area wireless networks (WWANs). Unlike in traditional wired and wireless IP-based networks, applications running over WWAN cellular environments are significantly affected by the vagaries of the cellular wireless medium. Prior research has proposed and analyzed optimizations at individual layers of the protocol stack. In contrast, we introduce the fi ...

Keywords: 3G, CDMA 2000, GPRS, HTTP, TCP, UMTS, cellular, cross-layer interactions, multi-layer performance optimizations, proxy

10 Engineering web cache consistency



Jian Yin, Lorenzo Alvisi, Mike Dahlin, Arun Iyengar

August 2002 **ACM Transactions on Internet Technology (TOIT)**, Volume 2 Issue 3

Publisher: ACM Press

Full text available:  [pdf\(403.96 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Server-driven consistency protocols can reduce read latency and improve data freshness for a given network and server overhead, compared to the traditional consistency protocols that rely on client polling. Server-driven consistency protocols appear particularly attractive for large-scale dynamic Web workloads because dynamically generated data can

change rapidly and unpredictably. However, there have been few reports on engineering server-driven consistency for such workloads. This article repo ...

Keywords: Cache coherence, cache consistency, dynamic content, lease, scalability, volume

11 Multicast Video-on-Demand services



Huadong Ma, Kang G. Shin

January 2002 **ACM SIGCOMM Computer Communication Review**, Volume 32 Issue 1

Publisher: ACM Press

Full text available: [pdf\(1.28 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The server's storage I/O and network I/O bandwidths are the main bottleneck of VoD service. Multicast offers an efficient means of distributing a video program to multiple clients, thus greatly improving the VoD performance. However, there are many problems to overcome before development of multicast VoD systems. This paper critically evaluates and discusses the recent progress in developing multicast VoD systems. We first present the concept and architecture of multicast VoD, and then introduce ...

Keywords: Quality-of-Service (QoS), VCR-like interactivity, Video-on-Demand (VoD), multicast, scheduling

12 High-capacity Internet middleware: Internet caching system architectural overview



Gary Tomlinson, Drew Major, Ron Lee

March 2000 **ACM SIGMETRICS Performance Evaluation Review**, Volume 27 Issue 4

Publisher: ACM Press

Full text available: [pdf\(571.68 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [index terms](#)

Previous studies measuring the performance of general-purpose operating systems running large-scale Internet server applications, such as proxy caches, have identified design deficiencies that contribute to lower than expected performance and scalability. This paper introduces a high-capacity proxy cache service built upon a specialized operating system designed to efficiently support large-scale Internet middleware. It suggests that specialized operating systems can better meet the needs of the ...

13 Protocol service decomposition for high-performance networking



Chris Maeda, Brian N. Bershad

December 1993 **ACM SIGOPS Operating Systems Review , Proceedings of the fourteenth ACM symposium on Operating systems principles SOSP '93**, Volume 27 Issue 5

Publisher: ACM Press

Full text available: [pdf\(1.22 MB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In this paper we describe a new approach to implementing network protocols that enables them to have high performance and high flexibility, while retaining complete conformity to existing application programming interfaces. The key insight behind our work is that an application's interface to the network is distinct and separable from its interface to the operating system. We have separated these interfaces for two protocol implementations, TCP/IP and UDP/IP, running on the Mach 3.0 operating sy ...

14 Rover: a toolkit for mobile information access



A. D. Joseph, A. F. de Lespinasse, J. A. Tauber, D. K. Gifford, M. F. Kaashoek

December 1995 **ACM SIGOPS Operating Systems Review , Proceedings of the fifteenth**

**ACM symposium on Operating systems principles SOSP '95**, Volume 29

Issue 5

Publisher: ACM PressFull text available: [pdf\(2.18 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)**15** Higher-order distributed objects

Henry Cejtin, Suresh Jagannathan, Richard Kelsey

September 1995 **ACM Transactions on Programming Languages and Systems (TOPLAS)**, Volume 17 Issue 5**Publisher:** ACM PressFull text available: [pdf\(2.33 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

We describe a distributed implementation of Scheme that permits efficient transmission of higher-order objects such as closures and continuations. The integration of distributed communication facilities within a higher-order programming language engenders a number of new abstractions and paradigms for distributed computing. Among these are user-specified load-balancing and migration policies for threads, incrementally linked distributed computations, and parameterized client-server applicat ...

Keywords: concurrency, continuations, higher-order languages, message-passing

16 Network Simulation 1: Performance of a mixed shared/distributed memory parallel network simulator

Cameron Kiddle, Rob Simmonds, Brian Unger

May 2004 **Proceedings of the eighteenth workshop on Parallel and distributed simulation****Publisher:** ACM PressFull text available: [pdf\(154.92 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#)

Designing fast parallel discrete event simulation systems for shared-memory parallel computers is simplified by the efficient communication operations enabled by the common memory space. The difficulties involved in designing large shared-memory computers and the resulting high cost of even modest size systems has led to the proliferation of computer systems consisting of small shared-memory computers connected via low-latency message-passing interconnection networks. This paper describes how a n ...

Keywords: conservative parallel discrete event simulation, network simulation, scalable network simulation

17 Client-server computing in mobile environments

Jin Jing, Abdelsalam Sumi Helal, Ahmed Elmagarmid

June 1999 **ACM Computing Surveys (CSUR)**, Volume 31 Issue 2**Publisher:** ACM PressFull text available: [pdf\(233.31 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Recent advances in wireless data networking and portable information appliances have engendered a new paradigm of computing, called mobile computing, in which users carrying portable devices have access to data and information services regardless of their physical location or movement behavior. In the meantime, research addressing information access in mobile environments has proliferated. In this survey, we provide a concrete framework and categorization of the various way ...

Keywords: application adaptation, cache invalidation, caching, client/server, data dissemination, disconnected operation, mobile applications, mobile client/server, mobile computing, mobile data, mobility awareness, survey, system application

18 Process migration



September 2000 **ACM Computing Surveys (CSUR)**, Volume 32 Issue 3

Publisher: ACM Press

Full text available: pdf(1.24 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Process migration is the act of transferring a process between two machines. It enables dynamic load distribution, fault resilience, eased system administration, and data access locality. Despite these goals and ongoing research efforts, migration has not achieved widespread use. With the increasing deployment of distributed systems in general, and distributed operating systems in particular, process migration is again receiving more attention in both research and product development. As hi ...

Keywords: distributed operating systems, distributed systems, load distribution, process migration

19 Video Storage: Periodic broadcast and patching services: implementation, measurement, and analysis in an internet streaming video testbed



Michael K. Bradshaw, Bing Wang, Lixin Gao, Jim Kurose, Prashant Shenoy, Don Towsley, Subhabrata Sen

October 2001 **Proceedings of the ninth ACM international conference on Multimedia**

Publisher: ACM Press

Full text available: pdf(797.96 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Multimedia streaming applications can consume a significant amount of server and network resources. Periodic broadcast and patching are two approaches that use multicast transmission and client buffering in innovative ways to reduce server and network load, while at the same time allowing asynchronous access to multimedia streams by a large number of clients. Current research in this area has focussed primarily on the algorithmic aspects of these approaches, with evaluation performed via analysis ...

Keywords: patching, periodic broadcast, server

20 Improving Scalability of Network Emulation through Parallelism and Abstraction



Cameron Kiddle, Rob Simmonds, Brian Unger

April 2005 **Proceedings of the 38th annual Symposium on Simulation**

Publisher: IEEE Computer Society

Full text available: pdf(208.63 KB)

Additional Information: [full citation](#), [abstract](#)

One approach to network emulation involves simulating a virtual network with a real-time network simulator and providing an I/O interface that enables interaction between real hosts and the virtual network. This allows real protocols and applications to be tested in a controlled and repeatable environment. To reflect conditions of large networks such as the Internet it is important that the emulation environment be scalable. This paper examines improvements in scalability of the virtual network ...

Keywords: Scalable Network Emulation, Parallel Discrete Event Simulation, Simulation Abstraction, Fluid Simulation

Results 1 - 20 of 200

Result page: [1](#) [2](#) [3](#) [4](#) [5](#) [6](#) [7](#) [8](#) [9](#) [10](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2005 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)